High Resolution Parallel Particle-In-Cell Simulations of Beam Dynamics in the Spallation Neutron Source Linac

Ji Qiang, Robert D. Ryne, Barbara Blind, James H. Billen, Tarlochan Bhatia, Robert W. Garnett, George Neuschaefer and Harunori Takeda
Los Alamos National Laboratory
Los Alamos, NM 87545

Abstract

In this paper, we present results of using high performance parallel computers to simulate beam dynamics in an early design of the Spallation Neutron Source (SNS) linac. These are among the most detailed linac simulations ever performed. The simulations have been performed using up to 500 million macroparticles, which is close to the number of particles in the physical system. The simulations are fully three-dimensional, and utilize RF cavity field data from modeling over 400 RF cavities. Furthermore, they use an improved model of the beam dynamics within the cavities. Traditionally, small scale two-dimensional simulations have been performed on PCs or workstations. While such simulations are sufficient for rapid design and for predicting root mean square properties of the beam, large scale simulations are essential for modeling the tails of the beam. The large scale parallel simulation results presented here represent a three order of magnitude improvement in simulation capability, in terms of problem size and speed of execution, compared with typical two-dimensional serial simulations. In this paper we will show how large scale simulations can be used to predict the extent of the beam halo and facilitate design decisions related to the choice of beam pipe aperture.

I. INTRODUCTION

The design of next-generation accelerators will require a major advance in numerical modeling capability. For example, applications that utilize high intensity H^+ or H^- beams, such as waste transmutation systems and future spallation neutron sources, place extremely stringent requirements on the allowed beam loss. High resolution simulations and large scale visualization, coupled with analysis, are needed to explore the phenomenon of beam halo formation. Such simulations typically involve solving of the Poisson-Vlasov equations in six-dimensional phase space. The most common approach for solving the equations is the particle-in-cell (PIC) method because of its computational efficiency and reduced memory requirements compared with direct Vlasov solvers. The PIC method has been widely used to study the dynamics of high-intensity beams in accelerators [1–6]. On a serial computer, e.g. a PC or workstation, the computational time and machine memory associated with using a large number of numerical particles restricts that number, and limits the accuracy of the PIC calculation. Parallelism can significantly improve the accuracy of PIC simulations by allowing the use of a larger number of particles and a finer grid resolution. It can also dramatically reduce the computation time. The availability of parallel supercomputers is making it possible to perform beam dynamics simulations with unprecedented speed and resolution. For example, below we will show three-dimensional results using 100 million macroparticles that require ten hours on parallel computer. In contrast, a two-dimensional simulation of the same problem using one million macroparticles running on a PC takes approximately two days. This corresponds to a $500\times$ improvement in simulation capability in terms of problem size and speed of execution, accompanied by an improved physical model (i.e. a three-dimensional simulation instead of a two-dimensional one).

In this paper we will describe a new parallel PIC code for modeling intense beams in high intensity RF linacs. We will describe the algorithms used in the code and their implementation on parallel supercomputers. We will also show an example of this new capability by applying it to a large-scale simulation of the SNS linac.

The organization of this paper is as follows: the three-dimensional parallel PIC code is described in Section 2, the parallel beam dynamics simulations are discussed in Section 3, and the conclusions and future work are presented in Section 4.

II. PARALLEL PARTICLE-IN-CELL CODE

The primary code for performing parallel simulations of high intensity RF linacs is IMPACT (Integrated-Map and Particle Accelerator Tracking code). This code was developed under a DOE-supported Grand Challenge in computational accelerator physics [13]. The IMPACT code combines techniques of magnetic optics with those of parallel particle-in-cell simulation by using a split-operator method. In this approach, the Hamiltonian governing the motion of individual particles in the accelerator is separated into two pieces, $H = H_{ext} + H_{sc}$, where H_{ext} corresponds to externally applied fields and H_{sc} corresponds to space-charge fields. The effect of H_{ext} is treated by using map-based techniques of magnetic optics, while the effect of H_{sc} is treated by solving Poisson's equation on a 3-dimensional grid using an FFT-based convolution algorithm. The following describes the split-operator approach, the methods used to treat the external fields and the self fields, and their implementation on parallel computers.

Let M_{ext} denote the map corresponding to H_{ext} and let M_{sc} denote the map corresponding to H_{sc}. Then the map M corresponding to $H_{ext} + H_{sc}$, accurate through second order in the step size h, is given by,

$$M(h) = M_{ext}(h/2)\ M_{sc}(h)\ M_{ext}(h/2). \tag{1}$$

Each complete step involves the following: (1) transport of a numerical distribution of particles through a half step based on M_{ext}, (2) solving Poisson's equation based on the particle positions and performing a space-charge "kick" (i.e. an instantaneous change in momenta, since H_{sc} depends only on coordinates, hence M_{sc} only affects momenta), and (3) performing transport through the remaining half of the step based on M_{ext}. If the

space charge is intense, this algorithm can be applied repeatedly on successive pieces of a beamline element; if the space charge is weak, it may be possible to achieve good accuracy by computing the space charge kicks infrequently, in which case M_{ext} could correspond to a string of elements within a half-step. Thus, an important feature of this approach is that it enables one to use large time steps (i.e. large steps in the independent variable) in the regime of weak or moderate space charge. Essentially, it enables one to decouple the rapid variation of the externally applied fields from the more slowly varying space charge fields. If more accuracy is required, one can use the 4th order algorithm of Forest and Ruth [7] or a higher order algorithm using a method of Yoshida [8]. These methods require multiple space charge calculations per full-step as compared to a single calculation for the second-order method. The program execution time is dominated by the space charge calculation. Therefore, minimizing the number of times the space charge fields are calculated greatly reduces the execution time.

The current version of IMPACT uses linear maps to treat drift spaces, magnetic quadrupoles, and RF gaps. A unique feature of the code is the treatment of the RF gaps. Some linac codes such as PARMILA and LINAC use RF cavity data, interpolated from data at a few selected energies, to compute the coefficients of an approximate transfer map. Other codes perform fine-scale numerical integration of trajectories through externally applied fields. The philosophy used in IMPACT is the following: First, the gap transfer map only involves a small number of coefficients (12 in the linear approximation), so it can be computed with high accuracy and little computational cost compared with pushing $\sim$100M particles. Second, it is inefficient to perform fine-scale integration of trajectories, which would naively involve a large number of space-charge kicks, particularly when the motion due to external fields is dominated by a linear map. (This would be analogous to performing fine-scale integration of millions of trajectories in a system with a linear time-dependent force, when, alternatively, one could simply integrate the equations for the linear map and propagate particles by a simple matrix multiplication). The treatment of RF gaps in IMPACT involves integrating the equations of motion for the linear map, $dM/dz = JSM$,

where J is the matrix corresponding to the fundamental Poisson brackets, and where S is the matrix representation of the quadratic Hamiltonian corresponding to the RF gap in the linear approximation. This allows one to compute the gap transfer map including fringe field effects, and without making assumptions about the entrance energy of the synchronous particle. (The equations for the synchronous particle are computing "on the fly," along with the equations for the gap transfer maps.) If desired, this approach can be extended to higher order by using techniques for the computation of nonlinear transfer maps.

The treatment of the space charge in IMPACT, which is the most time-consuming part of the calculation, is based on a PIC approach and involves three steps: charge deposition onto a grid, field solution on the grid, and field interpolation from the grid to the particle positions. IMPACT uses a cloud-in-cell algorithm for charge deposition and field interpolation. The field solution is based on a convolution of the charge density and the Green's function of the potential. This three-dimensional convolution is performed using an FFT-based algorithm with zero padding to implement the open boundary conditions [11,12]. The motion is assumed to be nonrelativistic in the bunch frame, where the space-charge calculation is performed. The main relativistic effect, i.e. the reduction in the space-charge force due to the azimuthal magnetic field, is included by properly transforming the space-charge fields from the bunch frame to the lab frame.

Lastly, we mention the parallel implementation of the above algorithms. In a parallel particle simulation code running on a multi-processor computer, the macroparticle data and field data are distributed among all the processors. IMPACT uses a domain-decomposition approach, in which each processor is responsible for a part of the spatial domain. In order to minimize inter-processor communication and achieve high performance, it is essential that each macroparticle resides on the processor corresponding to the particle's spatial coordinates. As the macroparticles move out of the spatial domain of a processor during the simulation, they are passed to the other processors using explicit communication. This data movement, handled by a parallel *particle manager*, represents a distinguishing feature of parallel particle simulation codes compared with analogous serial codes [9]. In the IMPACT

code, a two-dimensional domain-decomposition approach is employed. The particle manager is implemented using a message passing programming paradigm [10].

In summary, some key features of the IMPACT code are the following:

- Use of canonical variables, (x, p_x, y, p_y, t, p_t), where t is time-of-flight with respect to the reference particle (equivalent to a phase when normalized), and p_t is the (negative) energy deviation with respect to the reference particle.
- Reference particle trajectory computed during the simulation.
- Beamline elements (drift spaces, magnetic quadrupoles, and RF accelerating gaps) treated in the linear approximation (i.e. transfer matrices), with the matrices computed during the simulation.
- The transfer matrix for rf accelerating gaps is found by integrating the equations of motion for the map using RF gap fields obtained from an electromagnetic solver.
- A space charge calculation using a 3D (x-y-z) Poisson solver with open boundary conditions.
- The electrostatic fields are calculated in the bunch frame and Lorentz transformed back to the lab frame.
- Relativistic effects, i.e. reduction in the longitudinal force due to the azimuthal magnetic field of the bunch, are included.

A detailed discussion of the IMPACT code can be found in Reference [13].

III. BEAM DYNAMICS SIMULATION IN SNS LINAC

The SNS linac is being designed to provide an average 2-mA-current H^- beam with 1001.5 MeV kinetic energy for delivery to a high energy beam transport line used to inject into a proton accumulator ring. The peak beam current is 56 mA at 402.5 MHz. Each

bunch contains 8.7×10^8 particles. This results in higher space-charge effects than the present world's highest average power proton linac at Los Alamos National Laboratory which contains about 5.3×10^8 particles per bunch. A schematic plot of the SNS linac configuration addressed in this paper is given in Figure 1 [14]. It consists of three types of RF structures: a drift-tube linac (DTL), a coupled-cavity-drift-tube linac (CCDTL), and a coupled-cavity linac (CCL). The DTL, operating at 402.5 MHz, accelerates the beam from 2.5 MeV to 20 MeV using a 4-$\beta\lambda$ FFDD focusing lattice period. This period is equivalent to 8-$\beta\lambda$ at 805 MHz. It provides strong transverse focusing and good acceleration efficiency in the low-β regime. The bore radius of the DTL starts at 1.25 cm and gradually ramps up to 1.5 cm after the first 41 RF gaps. The CCDTL, operating at 805 MHz, accelerates the beam from 20 MeV to 79 MeV, which provides good acceleration efficiency for intermediate β beams. Two RF segments with a synchronous phase of 90 degrees are used between the DTL and the CCDTL to provide further bunching and longitudinal matching of the beam. A quadrupole FODO lattice with a period of 12-$\beta\lambda$ provides strong transverse focusing in the CCDTL. The bore radius in the CCDTL is 1.5 cm. Each CCDTL RF segment contains 2 two-gap CCDTL cells. The CCL, also operating at 805 MHz, accelerates the beam from 79 MeV to 1001.5 MeV. The CCL has good acceleration efficiency for a high-β beam. A constant transverse FODO focusing period is maintained throughout the CCDTL and CCL to minimize the potential for beam-halo production due to mismatch. Two types of CCL RF segments are used in the SNS design studied here. One type of CCL segment contains 8-cell RF cavities. This type of segment has a 1.5 cm bore radius for the first 43 segments, and a 1.75 cm bore radius for another 49 segments. The other type of CCL segment contains 10-cell cavities. Here the bore radius is 2 cm to avoid possible beam loss at the high-energy end, which could cause serious radioactivation of the accelerating structure. Figure 2 shows the effective shunt impedances for the three types of RF structures. We see that a good acceleration efficiency can be obtained for the DTL at low energy, for the CCDTL at intermediate energy, and for the CCL at high energy. The detailed design of these RF structures is described in reference [15].

Parallel particle-in-cell codes provide the capability for high resolution beam dynamics simulations. Before describing our large scale simulation results, it is useful to compare the parallel IMPACT code against a design code, LINAC, which uses the same space charge subroutine (called SCHEFF) as the widely used design code PARMILA [6]. The LINAC code uses a modified $r-z$ Poisson solver to calculate the radial and longitudinal space charge forces. The space charge force in a transverse elliptically symmetric beam distribution is calculated by using a transformation to an equivalent azimuthally symmetric distribution. For the case of a transverse aspect ratio not far from one, it gives a good approximation of three-dimensional space charge effects. In our comparison of the IMPACT and LINAC codes, the same initial distribution of one million macroparticles was used at the beginning of the CCL for both simulation codes. IMPACT used a $64 \times 64 \times 64$ $x-y-z$ Poisson solver while LINAC used a 20×40 $r-z$ solver. Figure 3 shows the transverse and longitudinal rms size as a function of distance inside the CCL from the IMPACT code and the LINAC code. We see that both codes agree with each other very well. The transverse aspect ratio of the beam size varies between 1.0 and 1.5. In this regime, the SCHEFF subroutine in the LINAC code provides an excellent approximation of the three-dimensional space charge effects. Figure 4 shows the transverse and longitudinal emittance as a function of distance inside the CCL. Again, both codes are in good agreement. A difference of a few percent is observed in beam emittance when comparing the results from the codes. Since the emittance is a sensitive diagnostic, such differences are not surprising considering the differences in the Poisson solvers and the treatment of the external fields. In Figure 5, we give the transverse and longitudinal rms sizes as a function of distance using both the IMPACT and LINAC codes for an initially mismatched beam at the beginning of the CCL. Here, we have increased the x rms size by a factor of 2 in the x phase plane and decreased the y rms size by a factor of 2 in the y phase plane in the initial distribution. This gives an initial transverse beam size aspect ratio of about 4.0. For this mismatched case, we see that there exist significant differences between the prediction of the IMPACT code and the prediction of the LINAC code. These differences result from the $r-z$ approximation used inside the LINAC code, which is not

appropriate for a beam with a large transverse aspect ratio. Previous comparison of the SCHEFF subroutine to a 3D space charge subroutine also suggested that the $r-z$ approach used in SCHEFF, along with its correction factor to treat azimuthal asymmetry, gave good agreement up to but not beyond a ratio of three [16]. Therefore, a full three-dimensional PIC code is needed to accurately study beam dynamics under this circumstance.

The SNS design requires that the beam loss along the linac must be kept as low as $0.1-1\ nA/m$ to limit structure radioactivation and permit hands-on maintenance. This limit would be exceeded if roughly 2-15 particles per bunch were lost at a single location in the linac. A multiparticle simulation on a conventional computer, e.g. a PC, is limited to roughly 10^6 macroparticles for practical simulation purposes. Such a simulation can provide useful information about the rms sizes of the beam, but it cannot accurately predict the particle distribution in the outer fringes of the beam due to its limited numerical resolution. In particular, a beam bunch in the SNS contains 8.7×10^8 particles. Therefore, in a 1 million particle simulation, each macroparticle corresponds to 870 physical particles, far above the regime we are interested in. In contrast, multi-particle simulations on parallel computers can easily be run with 100 million macroparticles. Such large-scale simulations can provide quantitative information about the tails of the distribution and the maximum amplitude of the beam, which is an important factor for choosing the design aperture of the machine. Figure 6 shows the maximum amplitude of the beam as a function of energy from LINAC simulations on a PC using 1000, 10000, 100000, and 1 million macroparticles. We see that with increasing number of macroparticles, the maximum amplitude of the beam also keeps increasing. Therefore, more macroparticles are required to determine the limiting value of the maximum amplitude of the beam and to provide good resolution around the outer fringes of the beam. Figure 7 shows the maximum amplitude of the beam from three-dimensional parallel simulations using the IMPACT code with 100 million and 500 million particles. It is seen that the increase in maximum amplitude is approaching a limiting value above roughly 100 million macroparticles, though this number can vary with the details of the physical problem.

Our parallel beam dynamics simulation of the SNS linac starts at the beginning of the DTL. The code advances particles through drift spaces, quadrupole fields and RF gaps. The dynamics inside the gaps is computed using external fields calculated from the electromagnetic code SUPERFISH [17]. Figure 8 shows the three longitudinal RF fields on the axis in one cavity for three types of RF structures used in the SNS design, the DTL, the CCDTL, and the CCL. For each integration step involving an rf gap, the axial electric field and its derivative with respect to z are read into the IMPACT code and used to compute the associated transfer map. The amplitude of the field is determined from a design code like PARMILA or LINAC so that the correct energy gain is achieved [14]. The initial beam energy at the beginning of the DTL is 2.5 MeV and the initial synchronous phase is -45 degrees. The initial particle distribution is from a LINAC code simulation on a PC using $10,000$ macroparticles. In the parallel simulation, we have used 100 million macroparticles with a $128 \times 128 \times 128$ physical grid to reduce statistical fluctuations and improve the numerical resolution of the PIC simulation. To generate an initial particle distribution with 100 million macroparticles, the original particle distribution has to be repopulated. A single numerical particle in the original distribution is replaced by 10000 numerical particles sampled uniformly from a six-dimensional box having its center at the phase space position of the original particle. The sizes of the boxes are chosen to be large enough so that the particles in each box disperse during the simulation, but small enough so that the new particle distribution has nearly the same initial rms sizes and emittances as the original distribution. Machine imperfections are taken into account in the simulation. Here, we have included the quadrupole magnet and RF cavity position errors, quadrupole rotation errors, quadrupole focusing gradient errors, static segment field-amplitude errors, errors in the distance between end-gaps of adjacent segments, static/dynamic RF module field-amplitude errors, and static/dynamic RF module phase errors in the CCDTL and CCL. The limits of these errors are given in Table 1 and Table 2. The position errors of quadrupoles and RF cavities and the rotation errors of quadrupoles are handled by transforming the particles into a new coordinate system. This is done by using a shift and rotation at the beginning

of each map, M_{ext}, in a quadrupole or RF cavity, applying the external field map M_{ext}, and transforming the particles back to the original coordinate system using a rotation and shift after M_{ext}. Due to the position misalignment errors, the centroid of the beam bunch will no longer be on the accelerator axis. In a real machine, steering magnets and beam position monitors are used to steer the beam back to the axis. In the simulation, the steering is handled by forcing the beam's horizontal and vertical displacement back to the axis at the location of a beam position monitor in the horizontal and vertical directions. Here, the beam position monitor was assumed to be located in the middle of a magnet for every 11 quadrupoles. The quadrupole focusing gradient errors and the RF field amplitude and phase errors are handled by randomly sampling a uniform distribution with a given error limit. Figure 9 gives the maximum transverse displacement and rms size of the bunched beam in the SNS linac with one set of all the above described errors. We see that the maximum particle amplitude is below the aperture size of the linac. This margin is needed to operate the linac safely and to avoid beam loss at the high energy end. The jump in rms beam size between the DTL and CCDTL at 20 MeV is due to a change of focusing lattice period from $8\beta\lambda$ to $12\beta\lambda$ at 805 MHz. Figure 10 gives the maximum longitudinal phase displacement and rms phase width of the beam. We see the gradual decrease of the longitudinal phase width due to acceleration in the linac. Figure 11 gives the normalized transverse and longitudinal rms emittances as a function of energy. The transverse emittance at the output of the CCL is 0.32 π-*mm-mrad*, and the longitudinal emittance is 2.3 *MeV-degree*. The transverse rms emittance is below the 0.5 π-*mm-mrad* limit specified as a ring injection requirement. Figure 12 shows the phase-energy plot of the final distribution at the end of the simulation. We see that the energy spread at the end of the linac is between -1.75 MeV and 2.11 MeV. The maximum energy variation must be less than ± 2.5 MeV at the linac output to inject into the high energy beam transport system.

The simulations shown above were done on an SGI Origin 2000 parallel computer using 256 processors. It took about ten hours to do the simulation using 100 million particles with a $128 \times 128 \times 128$ numerical grid. Typically, it will take the 2-dimensional LINAC code two

days to finish one simulation on a PC using 1 million particles and a coarser grid. For a 100 million particle simulation, the same code would take about 200 days to complete making it impractical for accelerator design use. Clearly, high performance parallel computing can drastically reduce the computation time when large numbers of particles are necessary to achieve the required numerical resolution.

IV. CONCLUSIONS

In this paper we presented the methods, implementation, and application of a new capability to model beam dynamics in high intensity RF linacs using parallel supercomputers. Our approach makes use of split-operator methods for Hamiltonian systems to combine features of parallel PIC codes with those of magnetic optics codes. As an application, we showed how the parallel IMPACT code was used to perform three-dimensional simulations of the SNS linac with 100 million macroparticles in only 10 hours. This is far beyond the capability of two-dimensional legacy codes that run on PCs. Such PC codes, which use a modified $(r-z)$ Poisson solver, have proven to be accurate for treating beam bunches that do not have a large transverse aspect ratio. While these PC codes are widely used in the design of RF linacs, they are not well-suited to studying the outer reaches of the beam halo due to their limited number of macroparticles. Parallel computing has enabled a roughly three order-of-magnitude increase in simulation capability, in terms of problem size and speed of execution, while concurrently enabling more realistic simulations (e.g. three-dimensional space-charge models instead of two-dimensional models). Nevertheless, the results shown here should be augmented with safety factors and with guidance from experience at accelerator facilities. A future version of IMPACT will include effects not included in the present study, such as third-order effects associated with quadrupoles and RF gaps, and image-charge effects from the beam pipe boundaries. Some changes of the linac configuration will likely appear in the future design. However, the simulations we have done here show an important application of parallel high-performance computing in studying the dynamics of high intensity beams

in linear accelerators.

ACKNOWLEDGMENTS

We are grateful to K. Crandall, B. Hartline, R. Jameson, and T. Wangler for fruitful discussions. Partial funding for this work was provided by the U.S. Department of Energy (DOE), Office of Science, Division of Mathematical, Information, and Computational Sciences, through the Grand Challenge in Computational Accelerator Physics. Partial funding was also provided by the U.S. DOE Office of Science, Division of High Energy and Nuclear Physics, through the Los Alamos Accelerator Code Group. Computational resources were provided by the Advanced Computing Laboratory located at Los Alamos National Laboratory and the National Energy Research Scientific Computing Center located at Lawrence Berkeley National Laboratory.

REFERENCES

[1] R. Ryne and S. Habib, in: Computational Accelerator Physics, ed. J. J. Bisognano and A. A. Mondelli, AIP Conference Proceedings 391, Woodbury, New York (1997) p. 377.

[2] B. B. Godfrey, in Computer Applications in Plasma Science and Engineering, ed. by A. T. Drobot, Springer-Verlag, New York, 1991.

[3] M. E. Jones, B. E. Carlsten, M. J. Schmitt, C. A. Aldrich, and E. L. Lindman, Nucl. Instr. and Meth. in Phys. Res. A318 (1992) 323.

[4] A. Friedman, D. P. Grote and I. Haber, Phys. Fluids B 4 (1992) 2203.

[5] Computational Accelerator Physics, ed. by R. Ryne, AIP Conference Proceedings 297, Los Alamos, NM 1993.

[6] H. Takeda and J. H. Billen, "Recent Developments of the Accelerator Design Code PARMILA," Proc. XIX International Linac Conference, Chicago, IL, August 1998, pp.156-158.

[7] E. Forest and R. Ruth, Physica D43 (1990) 105.

[8] H. Yoshida, Phys. Lett. **A 150**, 262 (1990).

[9] V. K. Decyk, Computer Physics Communications **87**, 87 (1995).

[10] M. Snir, S. Otto, S. H. Lederman, D. Walker, J. Dongarra, MPI: The Complete Reference, vol. 1, The MIT Press, Cambridge, MA, 1998.

[11] R. W. Hockney and J. W. Eastwood, Computer Simulation Using Particles, Adam Hilger, New York, (1988).

[12] W. H. Press, B. P. Flannery, S. A. Teukolsky, W. T. Vetterling, Numerical Recipes in C, Cambridge University Press, New York, (1988).

[13] J. Qiang, R. Ryne, S. Habib, and V. Decyk, "An Object-Oriented Parallel-In-Cell Code

for Beam Dynamics Simulation in Linear Accelerators," submitted to Journal of Computational Physics, 1999.

[14] T. Bhatia, et. al., "Beam Dynamics Design for the 1-GeV 2-MW SNS Linac", LA-UR-99-3802, Los Alamos National Laboratory, 1999.

[15] J. H. Billen, H. Takeda and T. Bhatia, "Linac RF Structures for the Spallation Neutron Source", Proceedings of the 1999 particle Accelerator Conference, New York, 1999, p.3585.

[16] G. Neuschaefer, private communication, February, 2000.

[17] J. H. Billen and L. M. Young, "POISSON SUPERFISH", Los Alamos National Laboratory report LA-UR-96-1834 (revision January 8, 2000).

TABLES

TABLE I. Table 1: Error limits for components of the SNS DTL

description of error	**limit on error**
quadrupole-gradient error	±0.5%
RF-field phase error	±0.5 degree
RF-field amplitude error	±0.5%

TABLE II. Table 2: Error limits for components of the SNS CCDTL and CCL

description of error	**limit on error**
quadrupole transverse displacement	±0.0127 cm
quadrupole roll	±5.0 mrad
quadrupole-gradient error	±0.25%
error in distance between the adjacent segment	±0.0254 cm
RF segment transverse displacement	±0.0508 cm
module field-amplitude error (dynamic)	±0.5%
module phase error (dynamic)	±0.5 degree
module field-amplitude error (static)	±1.0%
module phase error (static)	±1.0 degree
segment field-amplitude error (static)	±1.0%

List of Figures:

1. A schematic plot of SNS linac configuration.

2. Effective shunt impedances for the RF structures corresponding to the design bore radii.

3. (a) X rms size, (b) Y rms size (c) Rms phase width as a function of distance from IMPACT and LINAC simulations.

4. (a) X rms emittance, (b) Y rms emittance, (c) longitudinal emittance as a function of distance from IMPACT and LINAC simulations.

5. (a) X rms size, (b) Y rms size (c) Rms phase width as a function of distance from IMPACT and LINAC simulations with an initial mismatch.

6. LINAC simulation of CCDTL/CCL transverse beam profile with different numerical particles as a function of energy.

7. IMPACT simulation of CCL transverse beam profile with different numerical particles as a function of energy.

8. Longitudinal on-axis electric field distribution in (a) DTL, (b) CCDTL and (c) CCL.

9. IMPACT simulation of transverse beam profile in SNS linac as a function of energy.

10. IMPACT simulation of longitudinal beam profile in SNS linac as a function of energy.

11. IMPACT simulation of transverse and longitudinal rms emittance in SNS linac as a function of energy.

12. The phase energy distribution at the end of CCL.

FIG. 1. A schematic plot of SNS linac configuration.

FIG. 2. Effective shunt impedances for the RF structures corresponding to the design bore radii.

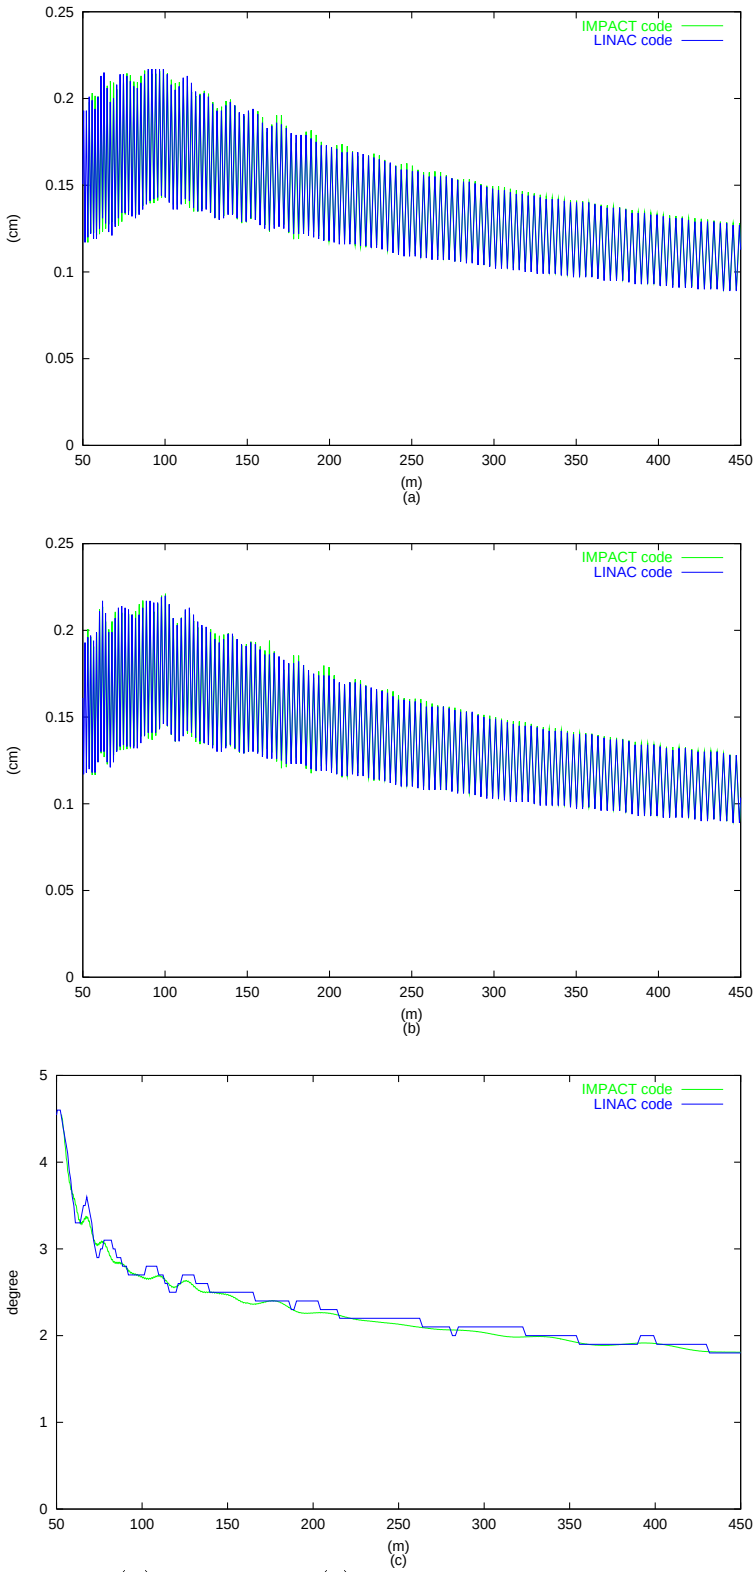

FIG. 3. (a) X rms size, (b) Y rms size (c) Rms phase width as a function of distance from IMPACT and LINAC simulations.

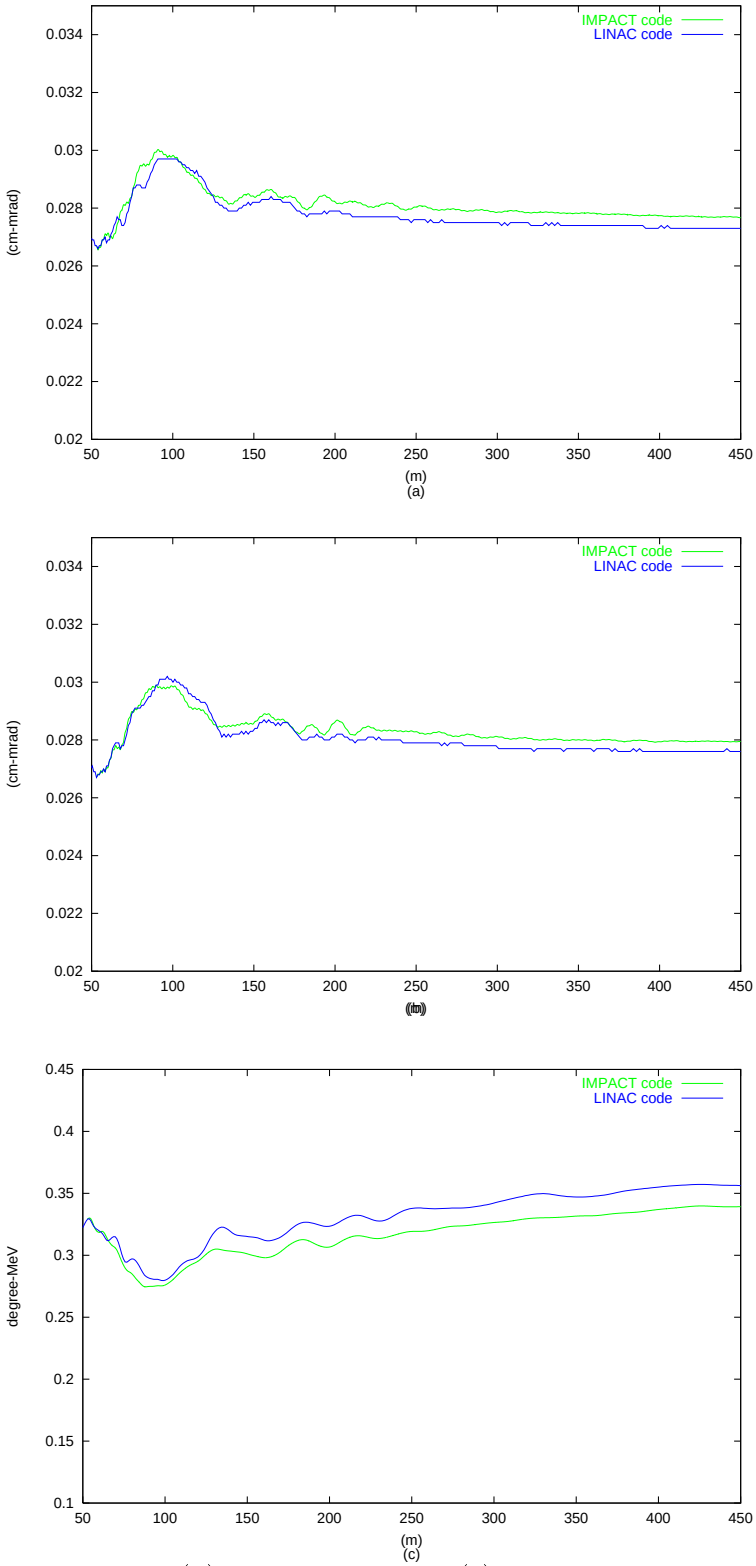

FIG. 4. (a) X rms emittance, (b) Y rms emittance, (c) longitudinal emittance as a function of distance from IMPACT and LINAC simulations.

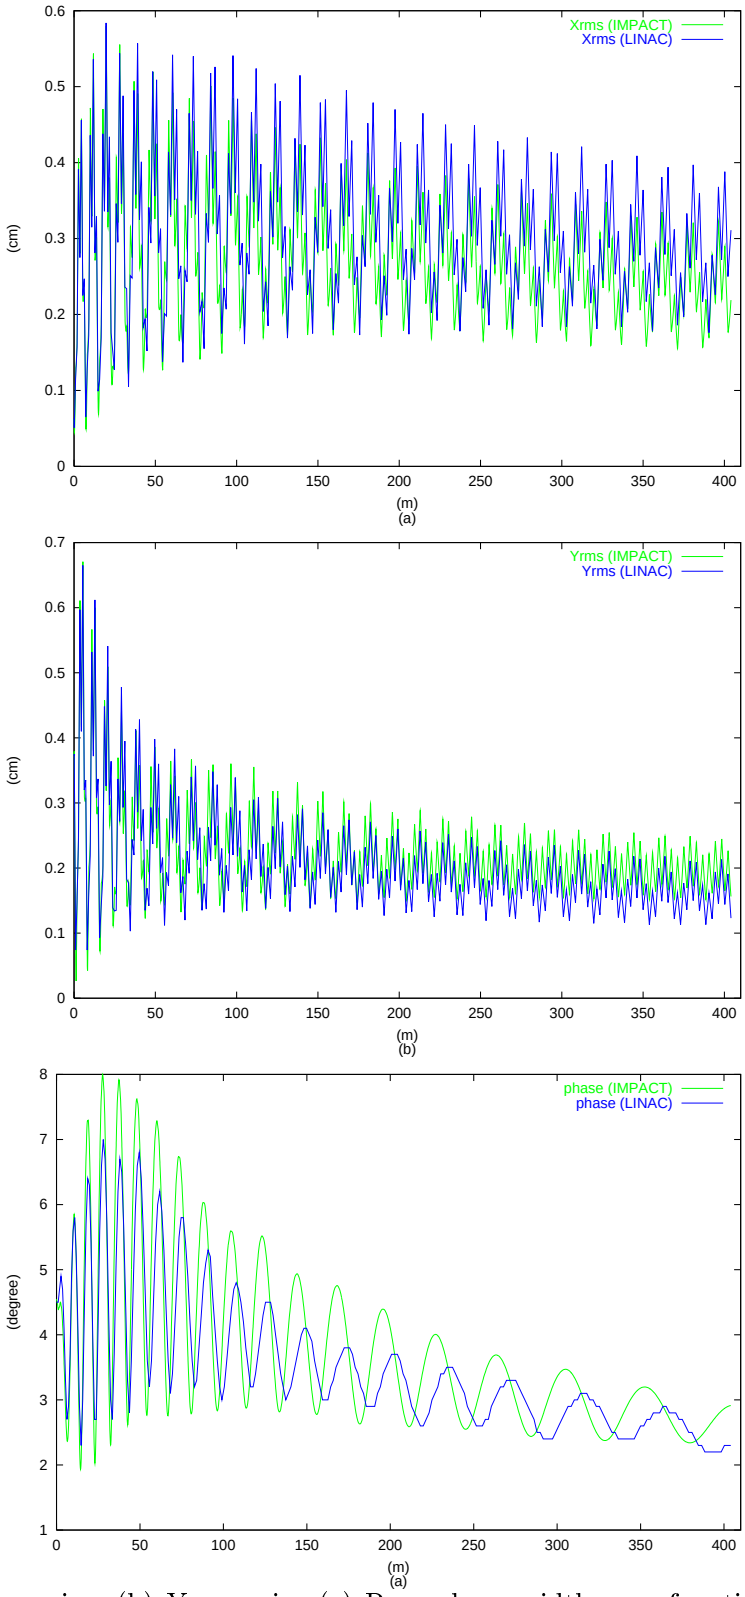

FIG. 5. (a) X rms size, (b) Y rms size (c) Rms phase width as a function of distance from IMPACT and LINAC simulations with an initial mismatch.

FIG. 6. LINAC simulation of CCDTL/CCL transverse beam profile with different numerical particles as a function of energy.

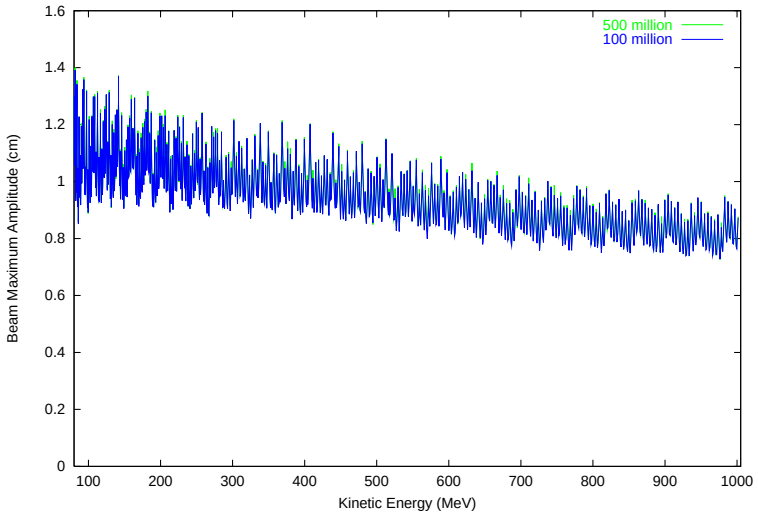

FIG. 7. IMPACT simulation of CCL transverse beam profile with different numerical particles as a function of energy.

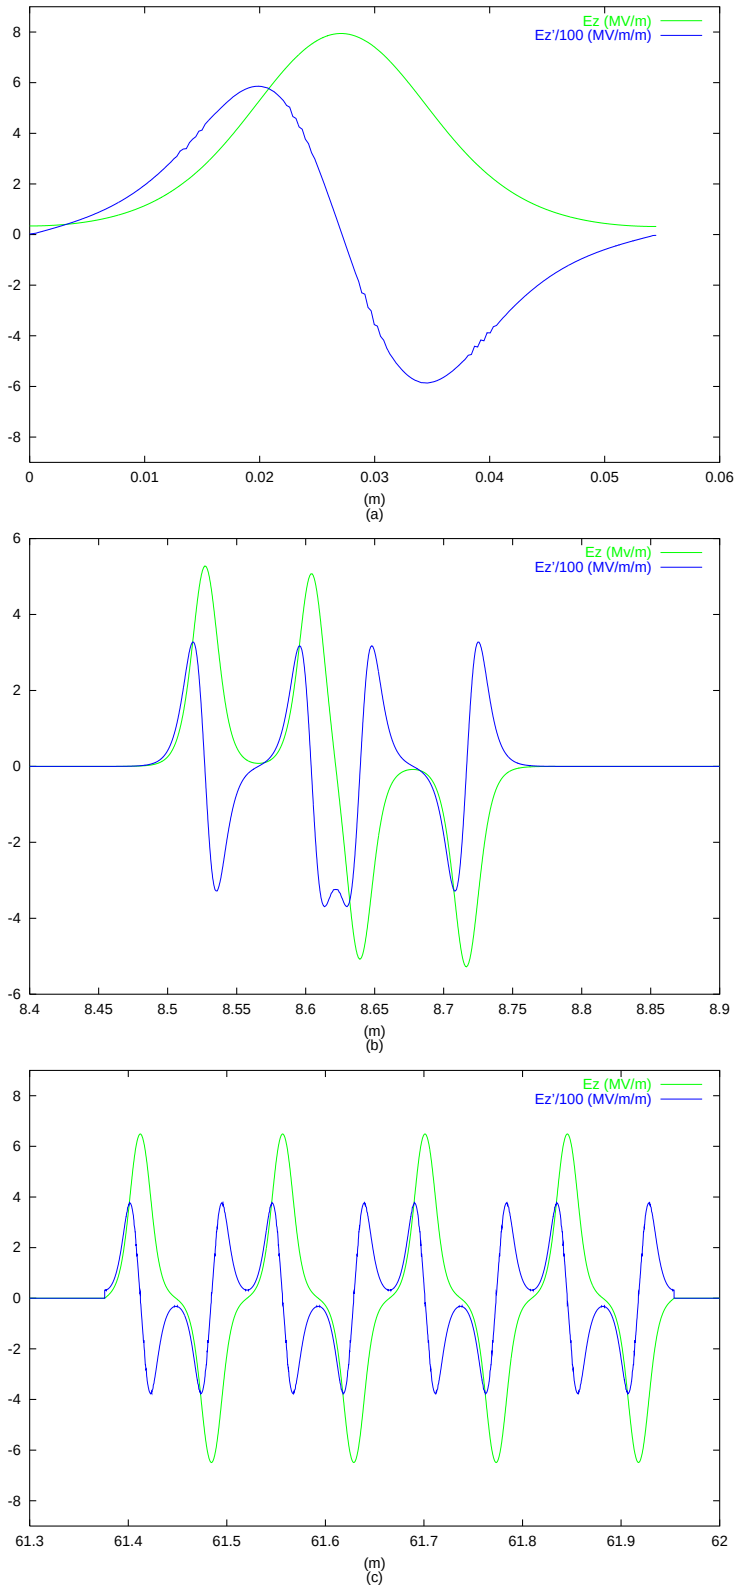

FIG. 8. Longitudinal on-axis electric field distribution in (a) DTL, (b) CCDTL and (c) CCL.

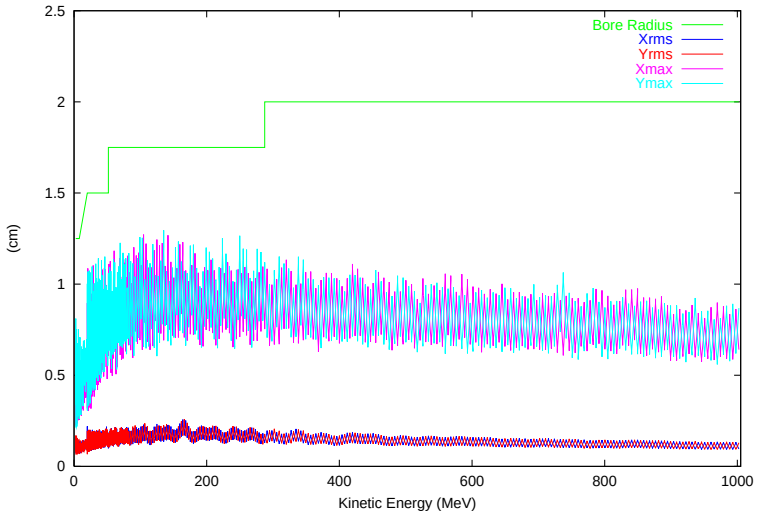

FIG. 9. IMPACT simulation of transverse beam profile in SNS linac as a function of energy.

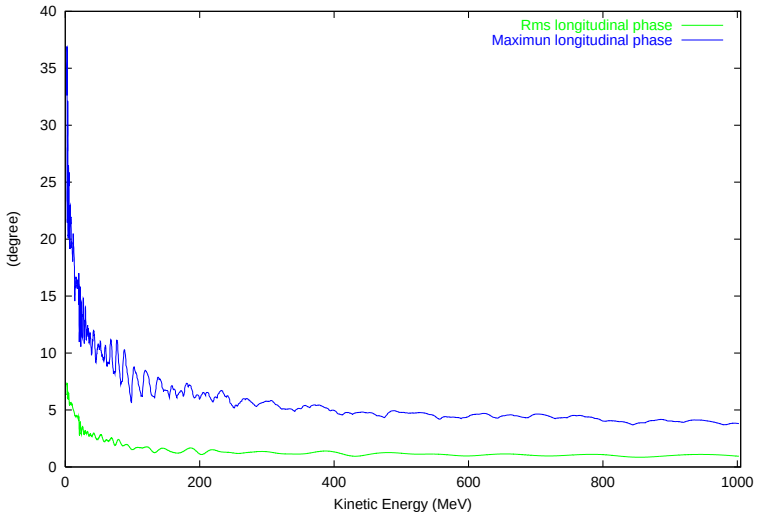

FIG. 10. IMPACT simulation of longitudinal beam profile in SNS linac as a function of energy.

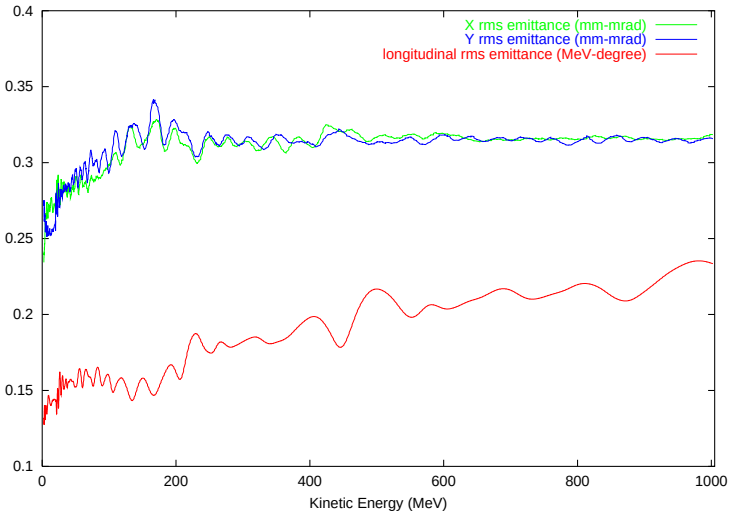

FIG. 11. IMPACT simulation of transverse and longitudinal rms emittance in SNS linac as a function of energy.

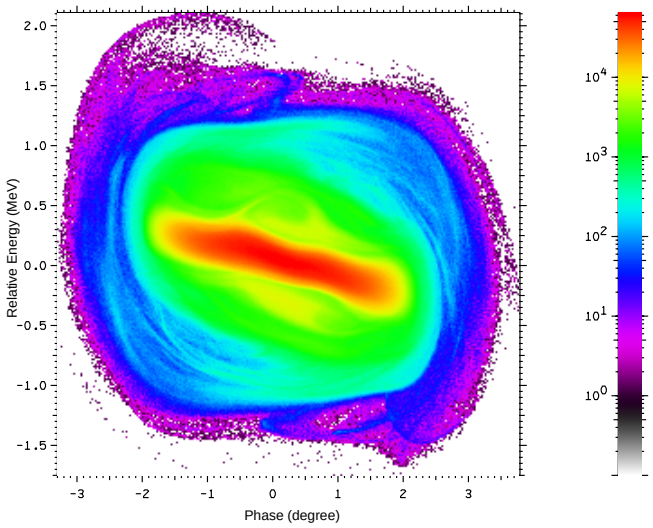

FIG. 12. The phase energy distribution at the end of CCL